


09/189,559

A70-892



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 1 148 417 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:

24.10.2001 Bulletin 2001/43

(51) Int Cl.7: G06F 11/22, G06F 11/14

(21) Application number: 01000104.8

(22) Date of filing: 05.04.2001

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR

Designated Extension States:

AL LT LV MK RO SI

(30) Priority: 20.04.2000 US 552862

(71) Applicant: International Business Machines
Corporation

Armonk, NY 10504 (US)

(72) Inventor: Snow, Paul Alan

c/o IBM United Kingdom Limited
Winchester Hampshire SO21 2JN (GB)

(74) Representative: Burt, Roger James, Dr. et al

IBM United Kingdom Limited,
Intellectual Property Law,
MP 110,
Hursley Park,
Hursley
Winchester, Hampshire SO21 2JN (GB)

(54) Application software damage detection and repair

(57) A method, system, and apparatus for detecting and repairing damaged portions of a computer system is provided. In a preferred embodiment of the present invention, a damage detection and repair facility monitors and detects changes to the computer system. The damage detection and repair facility compares these changes to the set of constraints defined by the working definitions for each application installed on the computer

system. The working definitions define the invariant portions of each application and define the constraints placed upon the computer system by each application. Responsive to changes that are in conflict with this set of constraints, the damage detection and repair facility makes such changes in the persistent storage so as to resolve these conflicts. This may be done, for example, by repairing a damaged file, installing a missing driver, or adjusting an environment variable.




FIG. 4B

EP 1 148 417 A2

Description

[0001] The present invention relates generally to the field of computer software and, more particularly, to methods of detecting and repairing damaged files and settings within a data processing system.

[0002] Currently, computer systems are built up through a series of installations provided by different software developers, each of which installs one or more different software components. There is no industry standard way to describe what comprises an application. Without this description, there is no way to implement a standard service that can protect the integrity of each application.

[0003] There is a need for a standard, simple, scalable, platform independent mechanism for detecting damaged applications and repairing them. Applications can be damaged in a variety of ways. Each of the following operations can damage one or more applications within a computer system:

- Installation of new applications
- Reconfiguring applications
- The use of an application
- Application error
- User error
- Viruses

Installation is dangerous because the current approach in the industry is a rather ad hoc approach in which the responsibility for the installation procedure for each application rests with each application's developer. During installation, any error or conflict between any two or more applications may potentially corrupt the computer system configuration. Applications do not have access to the requirements and dependencies of the other applications already installed on the target computer system. Without this information, no modification to a computer system's files and settings can be made completely safely. Yet the modification of files and settings is required to install applications onto computer systems.

[0004] The reconfiguring of an application may also require that changes be made to files and configuration settings needed or used by other applications, and such changes may render one or more of the other applications inoperable.

[0005] The use of applications may also require that the files and settings within a computer system be updated from time to time. This requires applications to be "well behaved" with respect to each other. Conflicts may still occur, either by chance, or error within an application. Yet the application developer's priority is always to their application without regard to potential conflicts with other applications, except to insure their application wins such conflicts. Because the requirements and constraints of each application are not defined, each developer also becomes responsible for supporting the configuration management of every system on which the

application is installed, and for handling conflicts in all configurations in which their application may be used. However, this responsibility is rarely, if ever, each application developer's top priority.

[0006] The current system of deploying and running applications defines only very weak barriers between applications to prevent a given application from damaging other applications. Any conflict, if handled in the given application's favor, may prevent a competitor's application from running. This state of affairs provides little motivation to avoid conflicts with a competitor's application so long as the developer's own application can be configured properly. Each application is dependent on its own installation and runtime code to verify its configuration and access to prerequisite devices, drivers, applications, and system settings.

[0007] Users themselves may cause problems by modifying files or configuration settings, either by accident, failure to follow a procedure properly, etc. Often all the files and configuration settings are exposed to modification by the user.

[0008] In reality, virus attacks account for a very small percentage of all application failures. Yet a viral attack can result in a huge amount of damage.

[0009] Currently the generally accepted method of protecting a computer system from damage from outside viral attacks is through the use of third party vendor virus protection software products. However, these products must be updated frequently to be able to detect the latest viruses. Any new virus that has been created since the software was updated may be undetected by the virus protection software, thus enabling that virus to corrupt or destroy files necessary for the proper performance of the computer.

[0010] Therefore, one needs to be able to strictly define an application's runtime representation. A Computer system, on the other hand, must allow modifications and changes in its persistent state in order to make the computer system useful. These requirements may be achieved by a method, system, apparatus, and data structure for strongly encapsulating an application and separating this representation from the runtime representation of the application. This encapsulated form is referred to as the Working Definition of the application. The defining characteristics of the application may be structured and managed within the application's Working Definition. The defining characteristics define the persisted state, settings, and structures required to build a valid representation of the application within its targeted computer system. This valid representation is referred to as the Runtime Representation of the application, and does not necessarily differ from what an installed application looks like within a computer system today.

[0011] One embodiment of such an arrangement has a memory for storing data for access by an application program being executed on a data processing system comprising: a data structure stored in said memory, said

data structure including: defining characteristics of an application program wherein the defining characteristics of the application comprise elements sufficient to construct a valid runtime representation of the application program within the data processing system.

[0012] There is, a need for a method, system and apparatus that automatically detects damaged files and applications and restore them to their proper condition.

[0013] According to one aspect of the present invention, a method of detecting and repairing damaged portions of a computer system is provided, the method comprising: detecting a change to the computer system; comparing the change to working definitions for each application installed on the computer system; wherein the working definitions comprise a set of constraints placed upon the computer system by each application installed on the computer system; and responsive to a determination that the change is in conflict with the set of constraints, modifying a persistent storage so as to resolve the conflict.

[0014] According to a second aspect of the present invention, a computer program product in computer readable media for use in a data processing system for detecting and repairing damaged portions of a computer system is provided, the computer program product comprising: first instructions for detecting a change to the computer system; second instructions for comparing the change to working definitions for each application installed on the computer system; wherein the working definitions comprise a set of constraints placed upon the computer system by each application installed on the computer system; and third instructions, responsive to a determination that the change is in conflict with the set of constraints, for modifying a persistent storage so as to resolve the conflict.

[0015] According to a third aspect of the present invention, a system for detecting and repairing damaged portions of a computer system is provided, the system comprising: means for detecting a change to the computer system; means for comparing the change to working definitions for each application installed on the computer system; wherein the working definitions comprise a set of constraints placed upon the computer system by each application installed on the computer system; and means, responsive to a determination that the change is in conflict with the set of constraints, for modifying a persistent storage so as to resolve the conflict.

[0016] The present application provides a method, system, and apparatus for detecting and repairing damaged portions of a computer system. In a preferred embodiment of the present invention, a damage detection and repair facility monitors and detects changes to the computer system. The damage detection and repair facility compares these changes to the set of constraints defined by the working definitions for each application installed on the computer system. The working definitions define the invariant portions of each application and define the constraints placed upon the computer

system by each application. Responsive to changes that are in conflict with this set of constraints, the damage detection and repair facility makes such changes in the persistent storage so as to resolve these conflicts. This may be done, for example, by repairing a damaged file, installing a missing driver, or adjusting an environment variable.

[0017] In carrying out the method, or utilising the computer program product or the computer system the correct version of the file may be retrieved from a server. This server may be accessed via a network and the network may be an Internet.

[0018] The novel features believed characteristic of the invention are set forth in the appended claims. For a better understanding of the invention, however, reference will now be made, by way of example, accompanying drawings, wherein:

Figure 1 depicts a block diagram illustrating a basic prior art computer architecture structure;

Figure 2 depicts a block diagram of a data processing system in which the present invention may be implemented;

Figure 3 depicts a block diagram of a personal digital assistant (PDA) in which the present invention may be implemented;

Figure 4A depicts a block diagram illustrating a data structure for strongly encapsulating an application in accordance with the present invention;

Figure 4B depicts a block diagram of a new model for a computer architecture in accordance with a preferred embodiment of the present invention;

Figure 5 depicts a portion of XML code demonstrating one method the requirements part of the working definition may be represented in accordance with a preferred embodiment of the present invention;

Figure 6 depicts a block diagram illustrating a method of automated damage detection and repair within a computing system in accordance with a preferred embodiment of the present invention;

Figure 7 depicts a flowchart illustrating an exemplary method of implementing a damage detection and repair facility in accordance with the present invention; and

Figure 8 depicts a flowchart illustrating an exemplary method of executing a strongly encapsulated application within a data processing system.

[0019] With reference now to the figures and, in par-

ticular, with reference to **Figure 1**, a block diagram illustrating a basic prior art computer architecture structure is depicted. Before proceeding, it should be noted that throughout this description, identical reference numbers refer to similar or identical features in the different Figures.

[0020] All existing computer architectures are patterned after a core, basic computer architecture **100**. This core, basic computer architecture is comprised of four elements: one or more processors **106**, one or more memory spaces **104**, one or more mechanisms for managing input/output **108**, and one or more mechanisms for defining persistence **102**. The Processor(s) **106** perform(s) the computations and instructions of a given application by loading its initial state into memory from the persisted image for that application.

[0021] Persistence **102** is the persisted storage of the applications apart from memory **104**. The basic nature of computer systems requires that all applications be stored somewhere. Even if a person types a program into a computer system every time they wish to use that program, the program is always stored somewhere, even if that storage is in someone's head. How an application is stored does not vary much from platform to platform. Applications are stored as, for example, executable files and libraries, files, databases, registry entries, and environment variables. Typically these files, databases, etc. are stored on a physical nonvolatile storage device such as, for example, Read only Memory chips, a hard disk, a tape, CD-ROM, or a DVD. Even in the most complex applications, the number of distinct persisted structures is really rather small, although there can be many of each type.

[0022] Regardless of the application's complexity, it depends on its persisted image. If the persisted structure is correct, the application will execute. If the persisted structure is wrong, the application will not execute.

[0023] With reference now to **Figure 2**, a block diagram of a data processing system in which the present invention may be implemented is illustrated. Data processing system **250** is an example of a computer which conforms to the architecture depicted in **Figure 1**. Data processing system **250** employs a peripheral component interconnect (PCI) local bus architecture. Although the depicted example employs a PCI bus, other bus architectures such as Micro Channel and ISA may be used. Processor **252** and main memory **254** are connected to PCI local bus **256** through PCI Bridge **258**. PCI Bridge **258** also may include an integrated memory controller and cache memory for processor **252**. Additional connections to PCI local bus **256** may be made through direct component interconnection or through add-in boards. In the depicted example, local area network (LAN) adapter **260**, SCSI host bus adapter **262**, and expansion bus interface **264** are connected to PCI local bus **256** by direct component connection. In contrast, audio adapter **266**, graphics adapter **268**, and audio/video adapter (A/V) **269** are connected to PCI local

bus **266** by add-in boards inserted into expansion slots. Expansion bus interface **264** provides a connection for a keyboard and mouse adapter **270**, modem **272**, and additional memory **274**. SCSI host bus adapter **262** provides a connection for hard disk drive **276**, tape drive **278**, and CD-ROM **280** in the depicted example. Typical PCI local bus implementations will support three or four PCI expansion slots or add-in connectors.

[0024] An operating system runs on processor **252** and is used to coordinate and provide control of various components within data processing system **250** in **Figure 2**. The operating system may be a commercially available operating system such as JavaOS For Business or OS/2, which are available from International Business Machines Corporation. Those of ordinary skill in the art will appreciate that the hardware in **Figure 2** may vary depending on the implementation. For example, other peripheral devices, such as optical disk drives and the like may be used in addition to or in place of the hardware depicted in **Figure 2**. The depicted example is not meant to imply architectural limitations with respect to the present invention. For example, the processes of the present invention may be applied to a multiprocessor data processing system.

[0025] Turning now to **Figure 3**, a block diagram of a personal digital assistant (PDA) is illustrated in which the present invention may be implemented. A PDA is a data processing system (i.e., a computer) which is small and portable. As with the computer depicted in **Figure 2**, and as with all computers, PDA **300** conforms to the computer architecture depicted in **Figure 1**. The PDA is typically a palmtop computer, such as, for example, a Palm VII[®], a product and registered trademark of 3Com Corporation in Santa Clara, California, which may be connected to a wireless communications network and which may provide voice, fax, e-mail, and/or other types of communication. The PDA **300** may perform other types of facilities to the user as well, such as, for example, provide a calendar and day planner. The PDA **300** may have one or more processors **302**, such as a microprocessor, a main memory **304**, a disk memory **306**, and an I/O **308** such as a mouse, keyboard, or pen-type input, and a screen or monitor. The PDA **300** may also have a wireless transceiver **310** connected to an antenna **312** configured to transmit and receive wireless communications. The processor **302**, memories **304**, **306**, I/O **308**, and transceiver are connected to a bus **314**. The bus transfers data, i.e., instructions and information, between each of the devices connected to it. The I/O **308** may permit faxes, e-mail, or optical images to be displayed on a monitor or printed out by a printer. The I/O **308** may be connected to a microphone **316** and a speaker **318** so that voice or sound information may be sent and received.

[0026] The computers depicted in **Figures 2** and **3** are examples of computers that conform to the computer architecture paradigm depicted in **Figure 1**. The sophistication of the system, number of components, and

speed of the processor may vary, but the basic model is the same as the architecture 100 depicted in Figure 1: means for input and output such as a keyboard and display (whether LED or a video display terminal), a processor, memory, and persistence. The persistence may be, for example, stored on a hard disk or hard wired into the system. However, the computers depicted in Figures 2 and 3 are merely examples. Other computers, in fact all current computers, also conform to this model. Examples of other such computers include, but are not limited to, main frame computers, work stations, laptop computers, and game consoles, both portable and conventional game consoles that connect to a television. Examples of game consoles include, for example, a Sony Playstation[®] and a Nintendo Gameboy[®]. Playstation is a trademarked product of Sony Corporation of Tokyo, Japan. Gameboy is a product and a registered trademark of Nintendo of America, Inc. of Redmond, Washington, a wholly owned subsidiary of Nintendo Co., Ltd., of Kyoto, Japan.

[0027] The computer of Figure 2 and the PDA of Figure 3 are also suitable to implement the data structures, methods, and apparatus of the present invention. The present invention modifies the computer architecture paradigm illustrated in Figure 1 in a way such as to present a new computer architecture paradigm with features and advantages as described below.

[0028] The present invention provides a data structure, process, system, and apparatus for allowing applications and computer systems to configure and manage themselves. The present invention also provides a process, system, and apparatus for creating encapsulated applications with controlled separation from an application's runtime representation. This separation is critical and necessary. The application's Encapsulation is provided by the Working Definition for the application, while the runtime image allows the application to execute within traditional execution environments, as defined by Windows, Unix, OS/2, and other operating system platforms.

[0029] With reference now to Figure 4A, a block diagram illustrating a data structure for strongly encapsulating an application is depicted in accordance with the present invention. In a preferred embodiment, any software application may consist of the following elements: identity 404, code 406, requirements 408, data 410, artifacts 412, and settings 414. These elements structure the application's defining characteristics. The defining characteristics define the persisted state, settings, and structures required to build a valid runtime representation of the application within its targeted computer system. The application's working definition documents and controls each of these elements of the application.

[0030] The application's identity 404 is defined by the application developer, and consists of the application's name, version, etc. The identity section 404 may also provide documentation, web links, etc. for the application useful both to automated management services and

users.

[0031] The code 406 represents the executable images, files, source code, and any other representation of the executable portion of the application that the application developer may decide to use. The code section 406 may also include executable images, files, source code, etc. to enable the installation, maintenance, configuration, uninstall, etc. of the application through the application's life cycle.

[0032] The requirements section 408 defines what directory structures, environment variables, registry settings, and other persisted structures must exist and in what form for this application to be properly constructed and installed in the persisted image of a computer system. The requirements section 408 also details any services and applications that are required to support this application. The requirements section 408 details these concepts as a set of constraints, and also may define the order in which these constraints must be resolved when such order is required. The requirements 408 also define the circumstances and constraints for the use of application specific installation, maintenance, configuration, uninstall, etc. code.

[0033] The Data section 410 defines data tables, configuration files, and other persisted structures for the application. The data section 410 is used to hold UI preferences, routing tables, and other information that makes the application more usable.

[0034] The Artifacts 412 are the persisted form of the value provided by the application. Examples of Artifacts 412 include documents, spreadsheets, databases, mail folders, text files, image files, web pages, etc. These are the files that contain the user's work ("user" in this context can be human or application).

[0035] The settings 414 are the persisted structures created or modified within the runtime representation that are intended to satisfy the requirements for the application. A distinction is drawn between the requirements 408 (which may specify a range of values in a registry setting or the form a directory structure must conform to) and the actual solution constructed in the runtime image 100 to satisfy these requirements.

[0036] The Requirements 408 are provided by the developers, application management software, environments, etc. Settings are made in an effort to resolve those requirements, and identify what requirements a setting is intended to satisfy.

[0037] With reference now to Figure 4B, a block diagram of a new model for a computer architecture 400 is depicted in accordance with a preferred embodiment of the present invention. The new computer architecture 400 of the present invention builds on the standard architecture as depicted in Figure 1. However, the new computer architecture 400 also includes an application's working definition 402. In a preferred embodiment, an application's Working definition 402 is an extensible markup language (XML) representation of the identity 404, code 406, requirements 408, data 410, artifacts

412, and settings 414; that is it includes all of the defining characteristics for the application. By defining these elements, separate from the runtime representation 100 of an application, working definitions 402 provides a way to strongly encapsulate an application which is compatible with operating systems that expect applications in an unencapsulated form. Working definitions define a new, universal picture of a computer system. Working definitions do not interfere with any operating system or platform because they have no active behavior. Working definitions define the application's runtime representation 100 by defining what an application is, what it requires, and what was done to give it what it needs, within a given computer system. While working definitions have no active behavior, they enable the implementation of a variety of automated services. These services include application life cycle management, resource tracking, application installation, damage detection and repair, computer system optimization, etc.

[0038] Strong encapsulation of state greatly reduces the global complexity of a computer system. States such as, for example, path requirements, file extensions, registry settings, program files, can be maintained in a pure state outside of the runtime representation 100. Because the runtime version of the application 100 is a redundant representation, it can be reconstructed, and fixed using the encapsulated version of the application as the standard. Important modifications can be persisted back to the application's Working Definition as a background task.

[0039] Working definitions use XML to define a flexible format for structuring the critical information about an application, thus encapsulating the application. XML provides the ability to extend working definitions to fit future requirements and, XML, just like the hypertext markup language (HTML) of the world wide web, can be as simple or as complex as the job requires. The actual size of a working definition is very small compared to the application it defines, even if the application is just a batch file. At the same time, working definitions can be used to describe vast, distributed, multi-platform applications.

[0040] Working definitions are platform and technology independent and address universal configuration issues. Furthermore, working definitions can be used as an open standard. Working definitions provide configuration information to services and to applications as well as providing bookkeeping support.

[0041] An application's valid runtime representation is one that satisfies a finite, structured set of constraints. Therefore, XML is an ideal method because XML provides an ideal representation for structured data and supports the ability to define links that allow tracking of internal relationships or for including references to structures defined externally. Furthermore, XML is designed to be easily extended and is platform independent.

[0042] However, XML is not a good representation for constant modification and access. The runtime repre-

sentation for the application provides the application with an efficient executable representation. To provide an application with an efficient executable representation, the XML can detail the required runtime structure and requirements for the application. Using these instructions, the application can be constructed and executed automatically. Prerequisites, in the proper order and with the proper configuration will also be constructed and executed automatically.

[0043] As the application executes, changes made to the application's state may be made as they are presently, that is, to the runtime representation. The XML specifies the important files to be persisted, and this persistence may be done in the background. The overhead is thus very minimal with little effect on the performance as viewed from a user's perspective.

[0044] The nature of the problem of automating management and configuration of computer systems requires a common, computer friendly and people friendly, complete representation of each software component. Without this information, the automated management of computer systems is not possible.

[0045] While the working definitions have been described herein and will be described henceforth with reference to an XML representation of the application, any number of other technologies and approaches might be used as well as will be obvious to one of ordinary skill in the art. For example, encapsulation could also be implemented using databases or with Zip files. Also, a database could be used for the actual implementation of the support for the encapsulation with XML used as a transfer protocol.

[0046] Furthermore, although described primarily with reference to constructing Working Definitions using XML, Working Definitions may be constructed using any structured format. Possible choices, for example, include the use of other structured forms. These include:

- Object based technologies (Java Beans, CORBA objects, C++ objects, etc.). These would have to be backed by some sort of database.
- Object Oriented databases. These are a mix of databases and Object based technologies.
- Functional technologies. One could build a Post-Script-like or Lisp-like language (compiled or interpreted) that defines these structures.
- Rule based technologies. Since Working Definitions generally resolve sets of constraints, this may well be the best way of implementing services that execute against the constraints that Working Definitions define. Working definitions could be constructed directly into Rule form, similar to forward chaining technologies such as ART, or backward chaining technologies such as Prolog.

- Tagged formats. XML is an example of a tagged format. However, there are other tagged formats that would work just as well. Tagged Image Format (TIFF), although generally used for defining images, may also be used to define Working Definitions since the use of custom tags is supported. Additionally, since TIFF is a binary format, it could hold the executable code and binary data. Other tagged formats include SGML, TeX, and LaTeX.

[0047] In addition to these examples of structured formats, there are certainly other formats and representations that would also work for constructing Working Definitions. Any technology that can 1) represent structured data, 2) support links to outside sources, and 3) can be used across platforms could be used to define Working Definitions.

[0048] With reference now to Figure 5, a portion of XML code is shown, demonstrating one method for representing the requirements element of the working definition, in accordance with a preferred embodiment of the present invention. The portion 500 of XML depicted is only snippet of a representation of the requirements element of the working definition and is given merely as an example. Furthermore, this portion 500 of XML is overly simplified, but does demonstrate how this information might be represented.

[0049] Platform tag at 502 reflects the requirement of the application that the platform be an xx86 type of computer, such as, for example, an Intel Pentium[®] class computer, running a Windows 95[®], Windows 98[®], or Windows NT[®] operating system. Pentium is a registered trademark of Intel Corporation of Santa Clara, California. Windows 95[®], Windows 98[®], and Windows NT[®] are all either trademarks or registered trademarks of Microsoft Corporation of Redmond, Washington.

[0050] Services tag at 504 indicates that the application requires the TCP/IP and file system services. Prerequisites tag at 506 indicates that the application Adobe Acrobat[®] is a prerequisite (i.e., requirement) for this application. Adobe Acrobat is a registered trademark of Adobe Systems Incorporated of San Jose, California.

[0051] The registry tag at 508 describes a registry entry. In this example, a Registry entry and a couple of keys are defined.

[0052] The Directory tag at 510 describes a couple of directories, the "programRoot" (which is required, but not specified by name) and a subdirectory "programRoot"/jclass.

[0053] Section 512 describes adding the programRoot directory to the path environment variable. The directory "programRoot"/jclass is also described as necessary to be added to the classpath environment variable.

[0054] The <test>...</test> tags in Section 512 describe tests that can be used to verify the setup of the enclosing section. This is one possible way of describing how to use application specific code to maintain an ap-

plication.

[0055] Figure 8, is a flowchart illustrating an exemplary method of executing a strongly encapsulated application within a data processing system is depicted in accordance with a preferred embodiment of the present invention. To begin, the strongly encapsulated application is transcribed into a runtime image (step 802). The defining characteristics embodied in the working definition of the encapsulated application are used in conjunction with system requirements of the platform on which the application is to be run are used in transcribing from the encapsulated version into the runtime version of the application.

[0056] Next, the runtime image of the application is translated into an executable version of the application by loading the runtime image into memory, thus producing an executable version of the application (step 804). The computer system's CPU then executes the instruction from the executable version of the application that has been loaded into the memory of the computer (step 808).

[0057] With reference now to Figure 6, a block diagram illustrating a method of automated damage detection and repair within a computing system is depicted in accordance with a preferred embodiment of the present invention. Currently, in the prior art, each application is responsible for its own integrity within a computer system. Automated facilities for detecting conflicts and resolving them, detecting damaged files or detecting missing registry entries are severely limited with prior art systems.

[0058] Conflicts and modifications to the runtime environment 100 cannot be avoided. They are the natural result of using a computer system, since the use of a computer system naturally leads to installing new applications, application updates, production and management of application artifacts, use of temporary files, and other changes in the system's configuration.

[0059] Damage detection and repair facility 602 uses the requirements defined in the working definitions 604 of all the installed applications as a set of constraints. Whenever changes occur to the runtime environment 100 or to any of the encapsulated applications, damage detection and repair facility 602 detects these changes and checks them against the set of constraints defined by the working definitions of the encapsulated applications installed on the computer system. The XML working definitions of each encapsulated application define the invariant portions of the application and the user's work which should be persisted as contrasted with temporary files that need not be persisted. An additional digitally signed section of the XML application working definition can provide insurance that checksums and files sizes for each of the application's files in the runtime environment 100 have not been modified. If a file has been modified, as detected by damage detection and repair facility 602, that file can be repaired with a signed, known valid version.

[0060] The set of files subject to inspection is very limited, since the working definitions define exactly which files should be checked. Instead of having to verify each and every file, the search can be limited to only those files the working definitions identify as critical. More exhaustive checks can be done, but the basic runtime representation verification does not require them.

[0061] With reference now to Figure 7, a flowchart illustrating an exemplary method of implementing a damage detection and repair facility, such as, for example damage detection and repair facility 602 in Figure 6, is depicted in accordance with the present invention.

[0062] The damage detection and repair facility monitors the computer system for changes (step 702) and determines whether the data processing system has received an indication to power down (step 703). If an indication to power down the data processing system has been received, then the process ends. If no indication to power down has been received, then the damage detection and repair facility determines whether changes have been made to any files, settings or encapsulated applications (step 704). Searches can be avoided when applications use a notification facility in order to facilitate this process. However, this kind of interaction with the application is not required. If no changes have been made, then the damage detection and repair facility 602 continues to monitor the computer system (step 702).

[0063] If changes have been made to some aspect of the runtime representation of an application as defined by that application's working definition, then the damage detection and repair facility compares the change against the constraints as defined by the set of working definitions in the computer system (step 706) and determines whether the change creates any conflict (step 708). If the change affects one or more working definitions with conflict, it is recorded in the settings section for the affected application(s) (step 710) and conflicts are resolved by restoring or adjusting the runtime representations affected. Changes to temporary files and settings (as defined by the working definitions of the applications that use them) do not cause conflicts.

[0064] In systems where security is very important, an optional test (step 711) for changes outside those defined as reasonable can be made. In this optional version, if the change does not create any conflict, then the damage detection and repair facility determines whether the change affects any working definition defined software component (step 711). This is done to detect those changes outside the defined constraints of the system. If the change does affect a working definition software component, then the settings of the affected working definitions are updated (step 712). If the change does not affect any working definition defined software component, then any number of strategies may be used (step 713) to address this security concern, including reversing the changes, logging the changes for inspection later, or reporting these changes to some monitoring facility. The damage detection and repair facility then con-

tinues to monitor the system (step 702). Thus, the working definitions of installed applications allow a computer system to be modelled as a set of applications that impose a set of constraints on the runtime representation of the computer system. These constraints define all of the elements of an application that can possibly be persisted as described above. When settings change, the damage detection and repair facility can evaluate these changes against this set of constraints as defined by the working definitions of each application. If the constraints are still met, then the settings can be recorded in the settings section of the affected working definitions. If the constraints are not met, the settings can be adjusted (and recorded) as required to restore the computer system back to its proper configuration.

[0065] It is important to note that while working definitions embody logically the totality of all of the defining characteristics of an application, the working definition may in fact be implemented as a distributed entity, rather than the single entity located on an individual machine as the present invention has been primarily described. Components of it may be accessed via links and may be physically stored on servers, hard drives, or other media and accessed over buses, the Internet, an intranet, or other channels.

[0066] It is important to note that while the present invention has been described in the context of a fully functioning data processing system, those of ordinary skill in the art will appreciate that the processes of the present invention are capable of being distributed in the form of a computer readable medium of instructions and a variety of forms and that the present invention applies equally regardless of the particular type of signal bearing media actually used to carry out the distribution. Examples of computer readable media include recordable-type media, such as a floppy disk, a hard disk drive, a RAM, CD-ROMs, DVD-ROMs, and transmission-type media, such as digital and analog communications links, wired or wireless communications links using transmission forms, such as, for example, radio frequency and light wave transmissions. The computer readable media may take the form of coded formats that are decoded for actual use in a particular data processing system.

[0067] The description of the present invention has been presented for purposes of illustration and description, and is not intended to be exhaustive or limited to the invention in the form disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art. The embodiment was chosen and described in order to best explain the principles of the invention, the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated.

Claims

1. A method of detecting and repairing damaged portions of a computer system, the method comprising:

detecting a change to the computer system;
 comparing the change to working definitions for each application installed on the computer system; wherein the working definitions comprise a set of constraints placed upon the computer system by each application installed on the computer system; and
 responsive to a determination that the change is in conflict with the set of constraints, modifying a persistent storage so as to resolve the conflict.

2. A computer program product in computer readable media for use in a data processing system for detecting and repairing damaged portions of a computer system, the computer program product comprising:

first instructions for detecting a change to the computer system;
 second instructions for comparing the change to working definitions for each application installed on the computer system; wherein the working definitions comprise a set of constraints placed upon the computer system by each application installed on the computer system; and
 third instructions, responsive to a determination that the change is in conflict with the set of constraints, for modifying a persistent storage so as to resolve the conflict.

3. A system for detecting and repairing damaged portions of a computer system, the system comprising:

means for detecting a change to the computer system;
 means for comparing the change to working definitions for each application installed on the computer system; wherein the working definitions comprise a set of constraints placed upon the computer system by each application installed on the computer system; and
 means, responsive to a determination that the change is in conflict with the set of constraints, for modifying a persistent storage so as to resolve the conflict.

4. A method according to claim 1 further comprising;
 or a computer program product according to claim 2 having instructions for;
 or a system according to claim 3 having means for:

responsive to a determination that the change

does not create a conflict with the set of constraints, updating settings in the working definitions of affected applications.

5. A method, computer program product or system according to one of the preceding claims, wherein the working definition of each application is provided by an extensible markup language representation.

6. A method, computer program product or system according to one of the preceding claims, wherein the step of modifying the persistent storage comprises repairing the runtime image of an application based on an encapsulated representation of the application.

7. A method, computer program product or system according to one of the preceding claims, wherein the step of modifying the persistent storage comprises repairing a damaged file using a correct version of the file from the working definition.

8. A method, computer program product or system according to one of the preceding claims, the method further comprising; the program having instructions for or the system having means for:

responsive to a determination that the change does not create a conflict with the set of constraints, determining whether the change affects any working definition defined software component; and
 responsive to a determination that the change does not affect any working definition defined software component, reversing the change.

9. A method, computer program product or system according to one of claims 1 to 7, the method further comprising; the program having instructions for or the system having means for:

responsive to a determination that the change does not create a conflict with the set of constraints, determining whether the change affects any working definition defined software component; and
 responsive to a determination that the change does not affect any working definition defined software component, logging the change.

10. A method, computer program product or system according to one of claims 1 to 7, the method further comprising; the program having instructions for or the system having means for:

responsive to a determination that the change does not create a conflict with the set of constraints, determining whether the change af-

fects any working definition defined software
component; and
responsive to a determination that the change
does not affect any working definition defined
software component, reporting the change to a 5
monitoring facility.

10

15

20

25

30

35

40

45

50

55




FIG. 1
(PRIOR ART)







FIG. 2





500

```

502 { <requirements>
      <platform>
      <processor>xx86</processor>
      <os>"Windows 95" or "Windows 98" or "Windows NT" </os>
      </platform>
504 { <services>
      <service>TCP/IP</service>
      <service>File System</service>
      </services>
506 { <prereqs>
      <application>Acrobat</application>
      </prereqs>
508 { <registryEntry name="HKEY_CURRENT_USER/Software/X.Com">
      <key>DebugPort<value>TCP/IP Port#</value></key>
      <key>Java VM<value>x<test>x gt "1.1.6"</test></value></key>
      </registryEntry>
510 { <directory name="programRoot"/>
      <directory name="jclass">programRoot"/jclass"</directory>
      <environmentVar name=path>
      <add>bin</add>
      </environmentVar>
512 { <environmentVar name=classpath>
      <add>jclass"/jappNetUI.jar"</add>
      <test>programRoot"/testClasspath"</test>
      </environmentVar>
      </requirements/>

```

FIG. 5

WORKING DEFINITIONS OF
INSTALLED APPLICATIONS

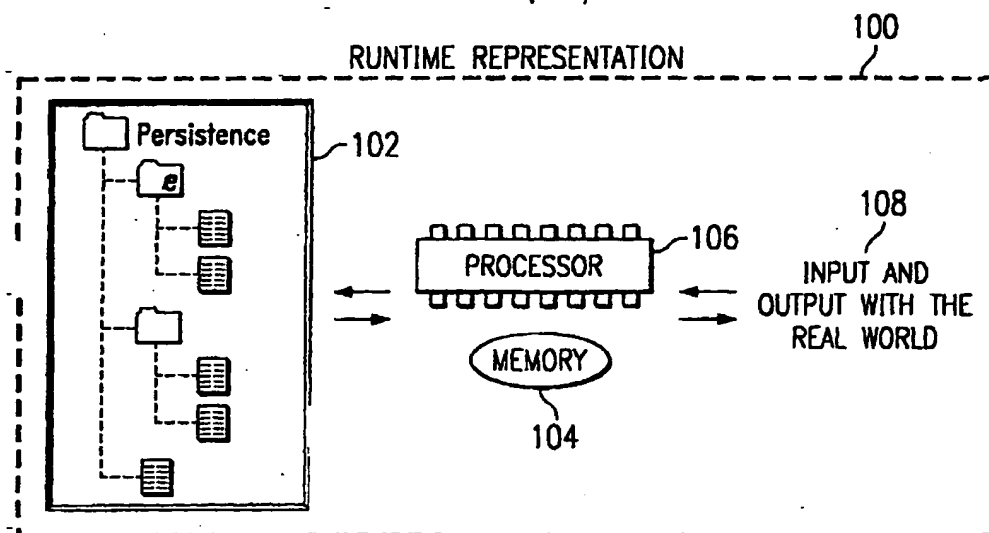
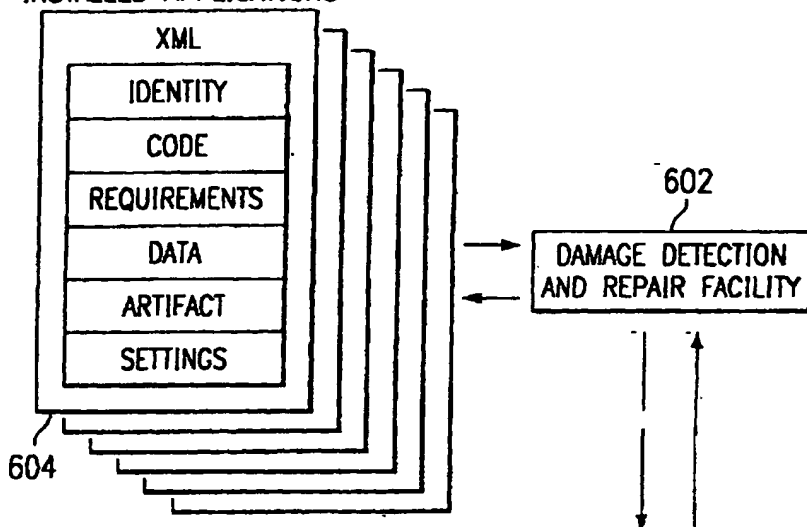
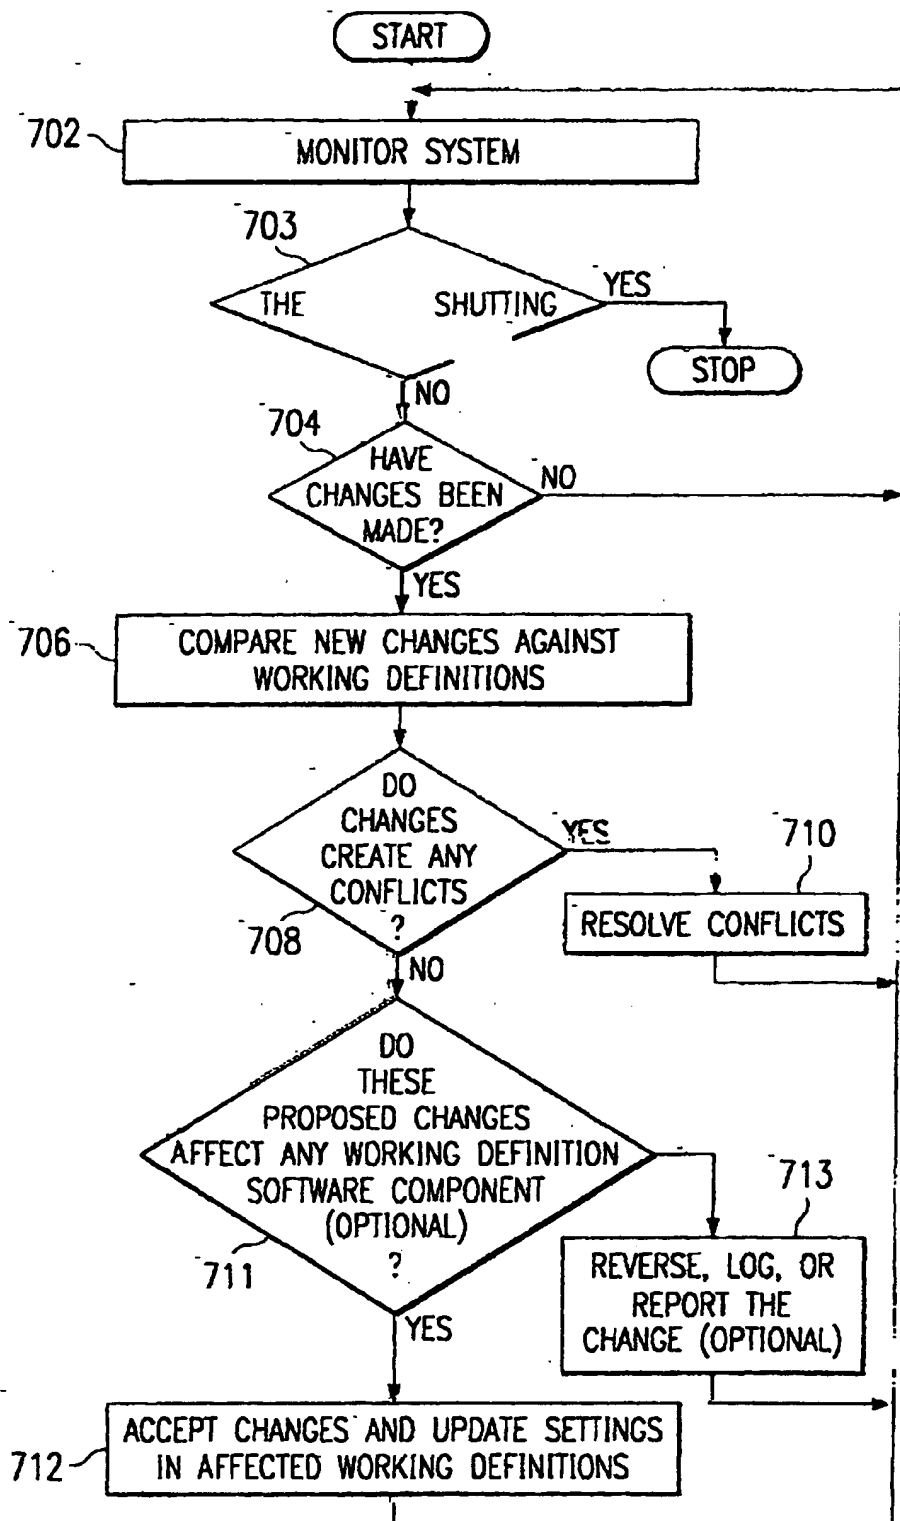


FIG. 6



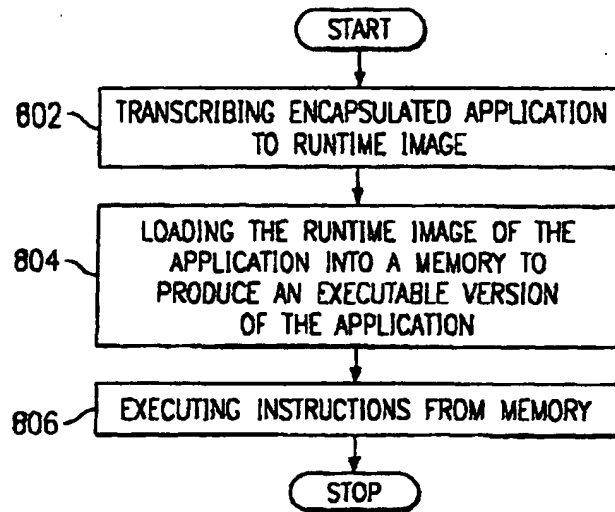


FIG. 8

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☒ **BLACK BORDERS**
- ☐ **IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- ☒ **FADED TEXT OR DRAWING**
- ☒ **BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- ☒ **SKEWED/SLANTED IMAGES**
- ☒ **COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- ☐ **GRAY SCALE DOCUMENTS**
- ☐ **LINES OR MARKS ON ORIGINAL DOCUMENT**
- ☐ **REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- ☐ **OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.